

# Chapter 1: Overview

## 1.1 About Control Room Logbook (CRL)

---

**Control Room Logbook (CRL)** is a multimedia, remotely-viewable, computer-based logbook for use in high energy physics experiment control rooms. **CRL** was conceived in early 1999. It was developed according to a set of requirements identified by physicists from several different experiments<sup>1</sup>. **Java** and **XML** were used to create this product in order to take advantage of the portability, ubiquitousness and rich functionality that these languages offer. **CRL** has three parts:

- the **CRL** application, installed in the control room to create, manipulate, and log entries
- the optional Process Logger which allows logging of entries created by external programs
- the optional **CRL** Web Access used to browse and annotate logged entries from virtually anywhere in the world

These three pieces are maintained and versioned independently for Linux, although they can still be installed and used together.

**CRL** supports a wide variety of entry data types. It can store the entries' content in both XML and HTML on the local disk drive or in a shared file system. In addition, it uses a relational database to store indexed information to query the entries. **CRL** is potentially compatible with any SQL-based relational database management system (RDBMS); currently it supports MySQL. The Web Access portion of **CRL** uses only the relational database and the HTML entries.

For information on implementing and configuring **CRL**, please see Part II: *CRL Administrator's Guide* and Part III: *CRL Administrator's Reference Manual*.

---

1. The experiments polled include DZero, CDF, CMS, NUMI, and BTeV.

## 1.2 The CRL Application Window

---

In this section, we provide an overview of the user interface, describing the “look and feel”, and identifying the types of objects found in the **CRL** window. The window in which **CRL** runs is composed of two frames (see image below):

**CRL toolbar**                      The **CRL** toolbar runs vertically down the left side of the window (see image below). This is a column of clickable buttons that control basic operations of **CRL**. It is not configurable, and therefore is the same in all installations.

**Desktop**                              To the right of the **CRL** toolbar, the desktop is the tabbed pane which takes up most of the window (see tabs at top of pane labelled **DETECTOR**, **DAQ**, and **ARCHIVE REPORT PAGE**). Each tab represents a different desktop page. The desktop displays one page at a time, according to the selected tab.

The desktop pages are configured by each experiment’s **CRL** administrator(s) according to the needs of the experiment.



## 1.2.1 The Look and Feel

**CRL** uses a familiar paradigm for its GUI, and includes features to aid in easy identification of desktop elements and logbook entry status:


- toolbars, buttons, cascading drop-down menus, and tooltips on mouse rollover
- separate desktop pages with tabbed views
- use of color for finding information or checking status “at-a-glance”
- automatic checkpointing

... and easy entry/manipulation of logbook data:

- drag-and-drop to initiate logbook entry
- double-click, control key and right-click shortcuts
- pop-up windows allowing user to type in or browse for item
- cut/copy/paste functions (within **CRL** and from/to external applications)
- form entry with option to save/reload data
- selection of individual entries or all entries
- “threads” for linking a series of related entries
- automatic, scheduled logbook data entry
- on-line help

## 1.2.2 The CRL Toolbar

The **CRL** toolbar is a column of clickable buttons that control basic operations of **CRL** independently of the desktop page operations:



<b>LOG IN/LOG OUT</b>	Brings up the login window (see section 2.3.2 <i>Logging In</i> )
<b>ENTRY SIGNERS</b>	Brings up a window in which to add/remove logged in users from the list of entry signers (see section 2.3.1 <i>About the Entry Signers Feature</i> )
<b>ADMINISTRATOR</b>	Brings up a window for adding/(de)activating/editing user information; administrative password required; (see Chapter 13: <i>Managing User Information</i> )
<b>CHANGE PASSWORD</b>	Brings up a window allowing users to change their password (see section 2.5 <i>Changing your Password</i> )
<b>SEARCH</b>	Brings up a window for querying database (see section 8.1 <i>Accessing Archived Entries from within CRL</i> )

**THREAD EXPLORER** Brings up a window for managing threads. Threads are used to link related logbook entries (see Chapter 7: *Threading Logbook Entries*).

**SCHEDULER** Brings up a window indicating all autoscheduled jobs currently running, and allows users to select any or each job for the purpose of changing scheduling parameters or terminating the job (see section 6.2.3 *Change Scheduling Parameters or Stop Job*).

### CURRENT SITUATION

Brings up a window for entering, changing or turning off a “current situation” message. A current situation message (also called a “global message”) gets automatically inserted in all subsequent entries until a user turns it off (see section 3.5 *Attaching Global Message to All Entries*).

## PREFERENCES

Brings up a window for selecting entry header lines to display/hide in container windows.



## CHECKPOINT RECOVERY

Entries not yet archived are checkpointed (saved) periodically to protect against loss. This button brings up a window listing all entries currently checkpointed for the machine (see section 4.5 *Recovering or Deleting Checkpointed Entries*).

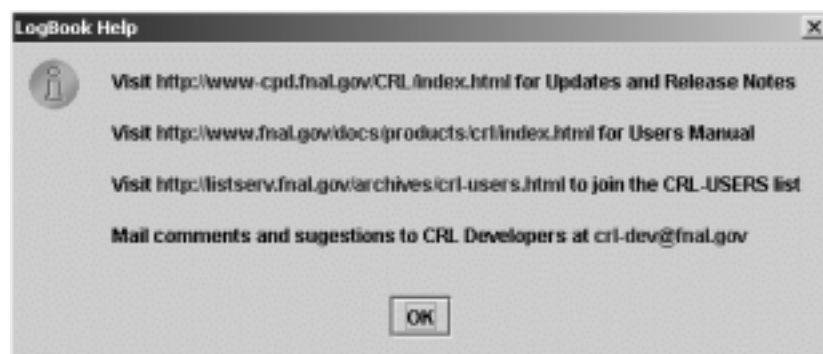
## ABOUT LOGBOOK

Brings up a window containing product release information, e.g.,



## HELP

Brings up a window showing the following information:



**EXIT LOGBOOK** Exits the **CRL** application (confirmation prompt provided)



### 1.2.3 The CRL Desktop

The large window to the right of the **CRL** Toolbar comprises the desktop. The **CRL** desktop is highly customizable, and each experiment configures the desktop pages, the input sources for logbook entries, and the hierarchical structure of categories and topics for logbook data storage according to its needs. The number of items needed on the desktop may therefore become quite large. Complex experiments with lots of monitoring equipment may choose to run **CRL** simultaneously on several computers in the control room, each configured to accept logbook entries from a particular set of sources, and to log the entries into a corresponding set of categories/topics. This allows the desktop on each machine to remain relatively uncluttered.

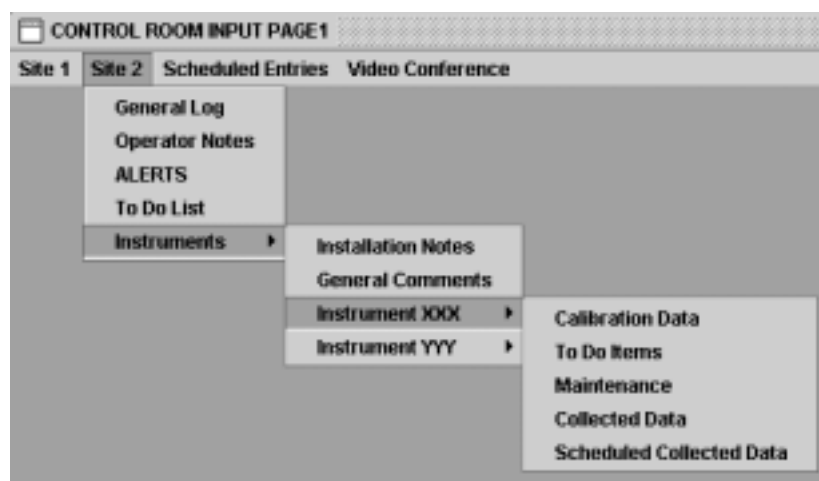
#### Objects on the Desktop

- |              |  |
|--------------|--|
| Page Tabs    | Page tabs are displayed horizontally along the top of the desktop. They identify the various desktop pages in your configuration. Click on one to make it the active page on the desktop.  |
| Desktop Page | A desktop page is the work space displayed for the selected tab. There may be several pages to your desktop; pages are configurable by experiment. A desktop page may be configured for data entry and manipulation, or for searching/viewing/manipulating archived entries only. All pages provide one or more menus, and each data entry page also provides a data-entry toolbar running vertically down the right-hand side of the page. Only one page is visible and active at a time. The active desktop page is identified at the top of the page (shown in this image as <b>TUTORIAL</b> , a data entry page), underneath the row of page tabs. |



## Pull-down Menus

Each desktop page has a set of menu headings lined up horizontally underneath the page title. Usually they represent general logbook entry categories. These are pull-down menus which have been configured by your experiment. They may cascade several levels in order to allow precise categorization of entries. The following image shows a desktop page labelled **CONTROL ROOM INPUT PAGE1**, and cascading menus starting from **SITE 2**. The menu options are described under *categories* and *topics* below.

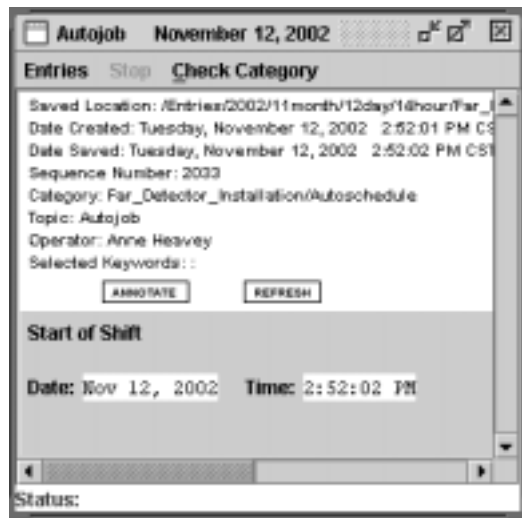



Categories	The menu headings and all the sublevels of categorization except the final one are considered logbook entry <i>categories</i> .
Topics	The final level of categorization is considered the <i>topic</i> (all menu options without an arrow pointing to the right are topics). Associated with each topic is a <i>container</i> , described below.
Containers	<p>A container is a <b>CRL</b> window that can contain one or more logbook entries. Each container represents and is labelled with a particular topic under the selected category hierarchy. You open a container of a given topic by double-clicking the menu item for that topic. Containers can be moved around the desktop, resized, iconized, and closed. They include a status bar at the bottom showing how many entries are selected within the container.</p> <p>There are several kinds of containers, described below, each with different properties. Each container provides menus for manipulating entries. These menus are not configurable by experiment, and they vary according to the container type.</p>




### *Input containers*

Input containers are containers in which you can add and edit new logbook entries. The input container at right contains entries that get stored under the experiment-configured topic **GENERAL INSTALLATION LOG** (notice the container's title).



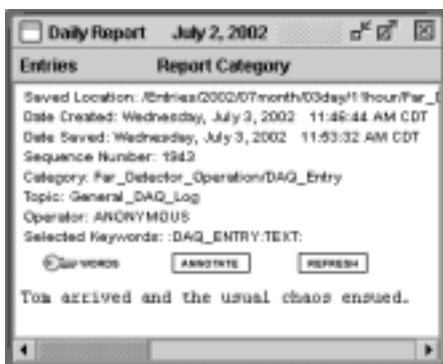

### *Scheduled Containers*

Scheduled containers are input containers for automatically scheduled logbook data entry (see Chapter 6: *Automated Logbook Entry and Archiving*). The container at left has been configured as **AUTOJOB**.




### Thread Containers

A *thread* links entries relating to a particular issue (see Chapter 7: *Threading Logbook Entries*). Thread containers are used for collecting threaded archived logbook entries, and are named by the user who creates the thread (e.g., **DAQ\_PROBLEM**).




### Report Containers

Report containers are used for collecting archived logbook entries to include in reports (see section 8.7 *Creating, Editing and Printing Reports in CRL*). They are configured by the experiment's **CRL** administrators. The one at left is named **DAILY REPORT**.



### Checkpoint Containers

Checkpoint containers are used for manipulating recovered checkpointed entries. They are configured by the experiment's **CRL** administrators. The one at right is named **CHECKPOINT**. Behind it we show the checkpointed entries.



## Logbook Entry Toolbar

On pages that allow logbook data entry, there is a toolbar running vertically down the right-hand side of the page. This toolbar includes a button for each logbook entry type (e.g., text, image binary file, **ROOT** data, forms, etc.), as configured for your installation. This toolbar is scrollable so that it can accommodate many data entry types.

To create an entry of a given type, you either double-click the button to create an entry in the currently selected open input container, or drag the button into any open input container.




## 1.3 Logbook Entries

---

### 1.3.1 Entry Format

Logbook entries get added to input containers or scheduled containers. In either type of input container, each data entry has a header, which identifies the date and time of the entry, the data category and topic, and the operator(s) who are logged in. Depending on the preferences set (see the **PREFERENCES** button under section 1.2.2 *The CRL Toolbar*), some header elements may not be displayed. If there are any keywords that can be attached to the entry (this depends on the configuration), then below the header you'll see a key symbol

(  WORDS ), which can be clicked to show or set keywords for the entry. The body of the entry is displayed under the key symbol. At the bottom, a status line indicates how many entries are currently selected.



Before a logbook entry is archived to the **CRL** database, its content and/or appearance can be modified, or it can be deleted. At this pre-archived stage, the data header appears in red. Once it is archived, the data entry can be annotated, but no longer modified or deleted, and the data header turns to black.

This color change was designed to let operators tell easily which entries are archived and which are not. In report containers, all headings are black, since all possible entries are already archived.

### 1.3.2 Entry Features

Data Types	Logbook data can be entered from a wide variety of sources. See section 3.3 <i>Logbook Data Entry by Type</i> .
Browse Directories	Several of the data types allow you to browse for a file or a URL to include in an entry. The system remembers the directory most recently browsed, and starts you from there each time. See sections 3.3.3 <i>Online Images</i> , 3.3.4 <i>Output Files from External Applications</i> .
Global Messages	A global message may be included in all logbook entries as experimental conditions warrant. See 3.5 <i>Attaching Global Message to All Entries</i> .
Editing	Before a logbook entry is archived to the <b>CRL</b> database, its content and/or appearance can be modified, or it can be deleted. See Chapter 4: <i>Editing Logbook Entries</i> .
Hot Key for Date/Time	The key sequence <b>CTRL-ALT-D</b> can be used for entering the current date and time into an entry or form field.
Double-click Shortcuts	

Double-click on toolbutton to create entry of corresponding type. Within container, double-click to create entry of type configured by your **CRL** admin.

## Control Key and Function Key Shortcuts

Short-cut key sequences are defined for several common operations.

Keywords	Keywords can be linked to logbook entries in order to provide an additional dimension for querying the database when attempting to later identify and retrieve particular entries. See section 4.4 <i>Attaching Keywords to a Logbook Entry</i> .
Checkpointing	Entries that have not yet been archived or deleted get written out periodically to a checkpoint directory. This protects against significant data loss in the event of a crash or accidental deletion. See section 4.5 <i>Recovering or Deleting Checkpointed Entries</i> .
Archiving	Once a logbook entry has been archived (saved), the entry gets “set in stone”. See Chapter 5: <i>Archiving Logbook Entries to the CRL Database</i> .
Reminders <sup>1</sup>	If an entry is left too long without being archived, your system may send messages periodically to ask you to address it. See section 5.4 <i>Entries Left Unarchived</i> .
Threading	A thread links a series of archived entries relating to a particular issue. It provides easy access to the flow of information on an issue, given that the entries may have been entered in different desktop pages, categories and/or topics, and by different operators. See Chapter 7: <i>Threading Logbook Entries</i> .
Searching/Retrieving	There are two methods of access to the database of archived entries from within <b>CRL: INQUIRIES</b> and <b>LOGENTRY EXPLORER</b> . See section 8.1 <i>Accessing Archived Entries from within CRL</i> . Entries can also be retrieved via the web, see 8.4 <i>Accessing Archived Entries on the Web</i> .
Annotating	The only change you can make to an archived entry is to add an annotation. This can be done from within <b>CRL</b> (see section 8.5.1 <i>Annotate Within CRL Application</i> ) and from a web browser (see 8.5.2 <i>Annotate from the Web Interface</i> ).

---

1. Also called "nagging".


## 1.4 Logbook Entry Storage

---

Up through version 1\_7\_06, **CRL** has been limited to storing/retrieving entries to/from a repository that is either local to the client machine or on a remote drive that is mounted on the client machine (thus appearing local to the user). As of version 1\_7\_07, **CRL** supports the capability to store entries in a remote repository accessed via a Remote File System Service (RFS). A repository of this type is called a *global file system*. Each client must be configured to send entries to a single repository only, but it now has the option to choose a global one, if available. A global repository can be accessed by clients either locally or remotely through the RFS Service.

The implementation of the Remote File System Service is a major enhancement to the product. It supports a number of features, including:

- distributed computing
- choice of destination file system (local or global)
- strong authentication (Kerberos V5 or other)
- mixing of Windows and Linux clients and servers



The RFS Service is the intermediary between the remote clients (of which there may be many) and the global file system (the repository). It enforces authentication and authorization for clients accessing the repository, and maintains each remote-client-specific configuration (each remote client may have its own view of the global repository).

As a **CRL** end user, you don't need to change the way you create, archive, and retrieve entries. The administrator configures the repository, and entries just go where they're supposed to go! Here is a summary of the start-up process for using the global repository (see the numbers 1, 2 and 3 in the diagram above):

- 1) Assuming Kerberos authentication is enabled in the RFS Service (as is required at Fermilab), the remote client must first authenticate to Kerberos. **CRL** then sends a request to the RFS Service to make available the client-specific view of the Global File System.
- 2) The RFS Service verifies authentication and authorization, and responds with the client-specific abstraction of the repository.

- 3) Client now has the abstraction of the repository mounted locally, and the user logs and retrieves entries to and from the (real, global) repository as though it were local. But behind the scenes, new entries are travelling to the RFS Service, being handled there as per configuration, and then being sent on to the global repository. Retrieval requests also go through the RFS Service.

